HAMMERSPACE

Activating Data for AI:

Unifying Data Without Rebuilding Infrastructure

Floyd Christofferson
VP Product Marketing -- Hammerspace

RMACC
Boise, May 13

Data underpins every AI solution…

… but it doesn't live in one place

73%
of Enterprise data is siloed

Today's Enterprise Needs…

Visibility into Data

Unified Access

Automated Pipelines

AI Exposes Limits of Centralized Data

AI is not just compute-constrained. It is data-constrained.

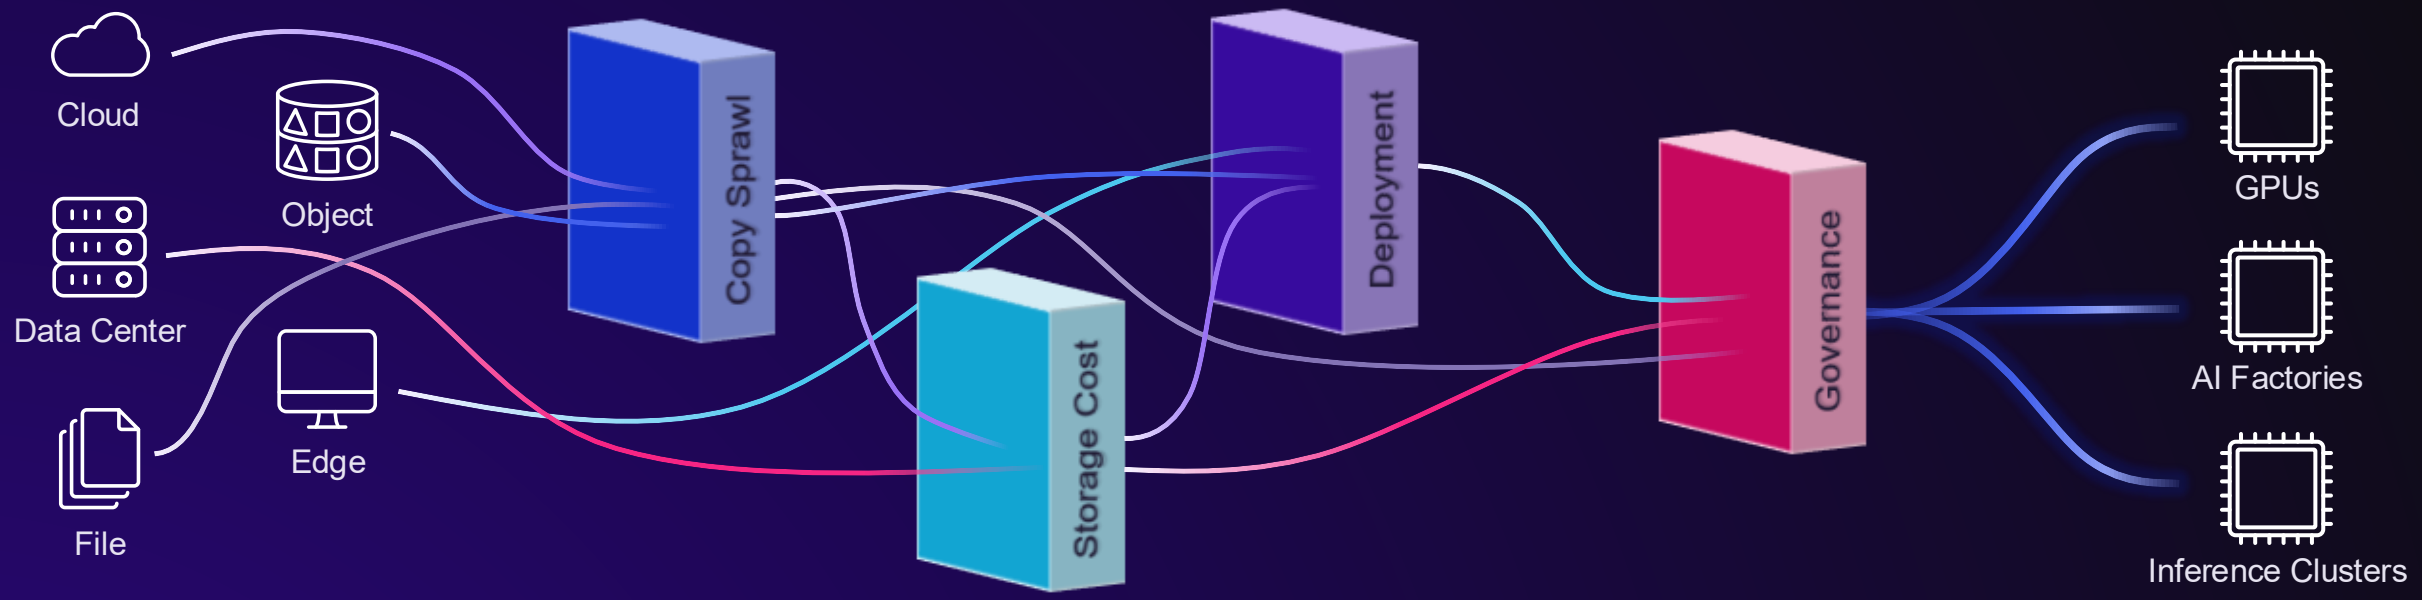

AI Changes How Data Must be Used	What Breaks At Scale
• Data stays where it is • Compute moves to where it's available • Workloads are no longer fixed	• Copy sprawl • Idle GPUs • Power & cost waste • Deployment delays

The AI FOMO Gap

- 57% of IT leaders say they were pushed to deploy before they were ready.
- Only 20% of AI investments are currently delivering ROI
- 14% of organizations are confident in their data governance

Enterprises Need an Outcome for AI Success

Not just a bunch of parts and tools to assemble themselves

Where should I host it? | **What tools do I use?** | **And how do I make sure it is efficient and secure?**

FOMO

(Fear of Missing Out)

- Board, Competition, and BU Pressure
- Those who don’t adopt will lose the market
- “Everyone’s doing it” pressure cycle

FOMU

(Fear of Messing up)

- Investing in the wrong technology
- Large capital outlay before return
- Will it deliver an ROI?
- Security and risk exposure

The Reality Has Changed

From HPC Built for Mod/Sim to AI + Hybrid Workflows

HPC Was Built for Mod/Sim

Compute-Intensive. Predicable. Centered.

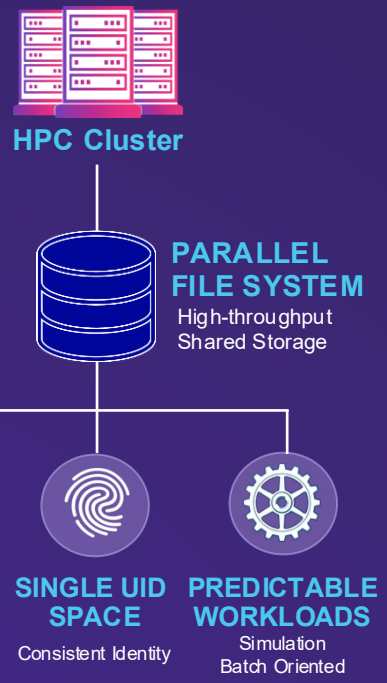

The Model That Worked

- ✓ Centralized storage
- ✓ Consistent identity
- ✓ Stable access patterns
- ✓ High-throughput shared file systems
- ✓ Built for long-running jobs

The Emerging Reality

Data-Intensive. Dynamic. Distributed.

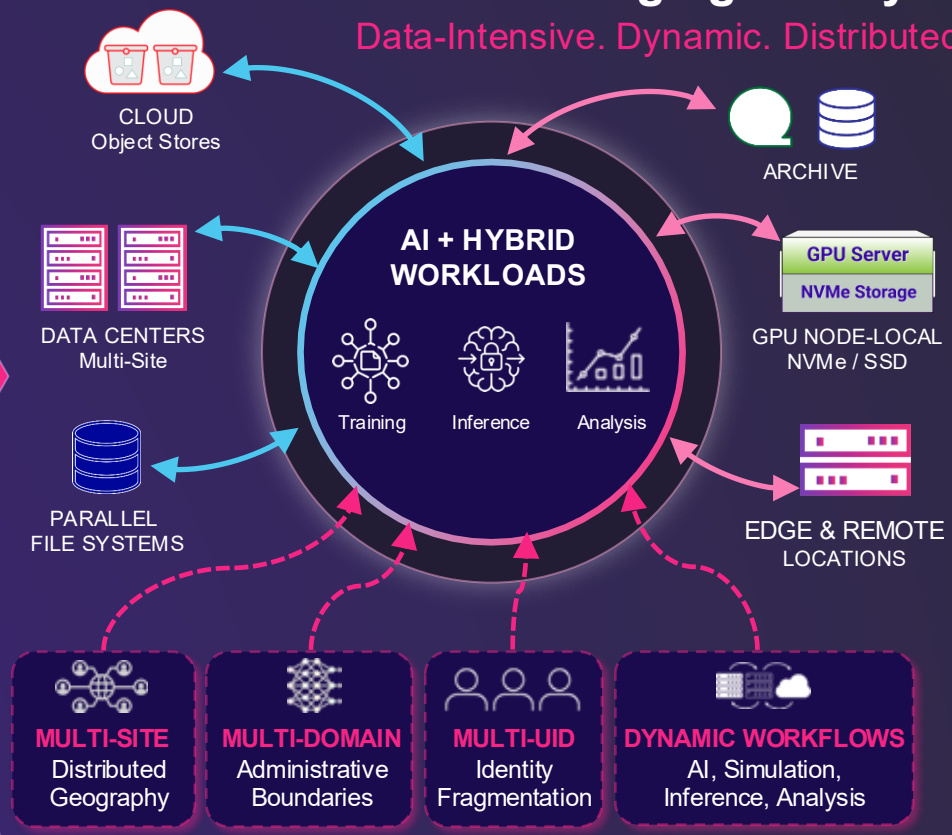

What's Different Now

- ✗ Data is everywhere, not just in one place
- ✗ Identities don't always match
- ✗ Workloads are dynamic and diverse
- ✗ GPUs need data fast and local
- ✗ Moving or copying creates friction, cost, and risk

This is not just a Performance Problem. It Is a Data Orchestration Problem

The Problem: Effectively Utilizing Siloed Data

Data Access is Fragmented

Complex for Users

- No global view or access
- Data copies explode, and workflows are siloed

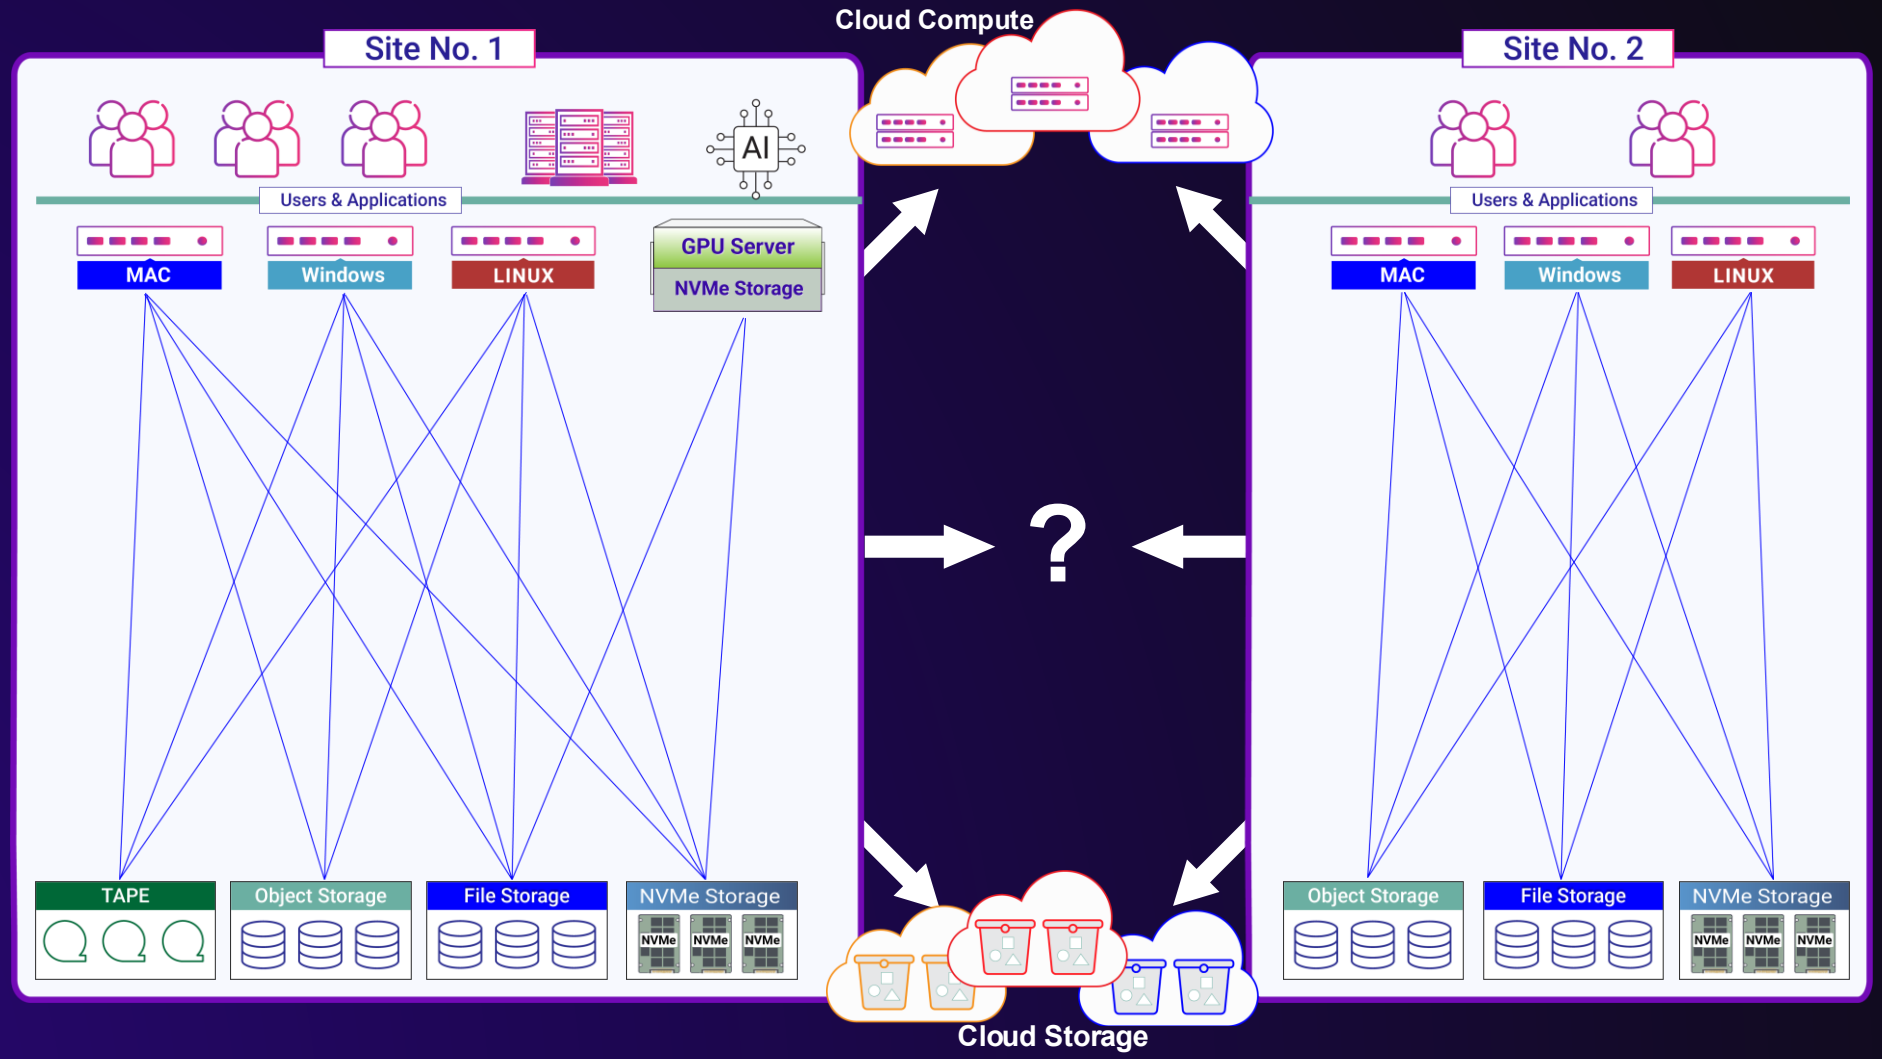

HAMMERSPACE

Multi-Vendor, Multi-Site, Multi-Cloud Storage

The Problem: Effectively Utilizing Siloed Data

Data Access is Fragmented

Complex for Users
- No global view or access
- Data copies explode, and workflows are siloed

Complex to Manage
- IT services are fragmented
- Data Protection limited by system/location, not global

Point Solutions Used to Fill Gaps

Slows Workflows
- High-performance is siloed
- Volume copies & migrations required

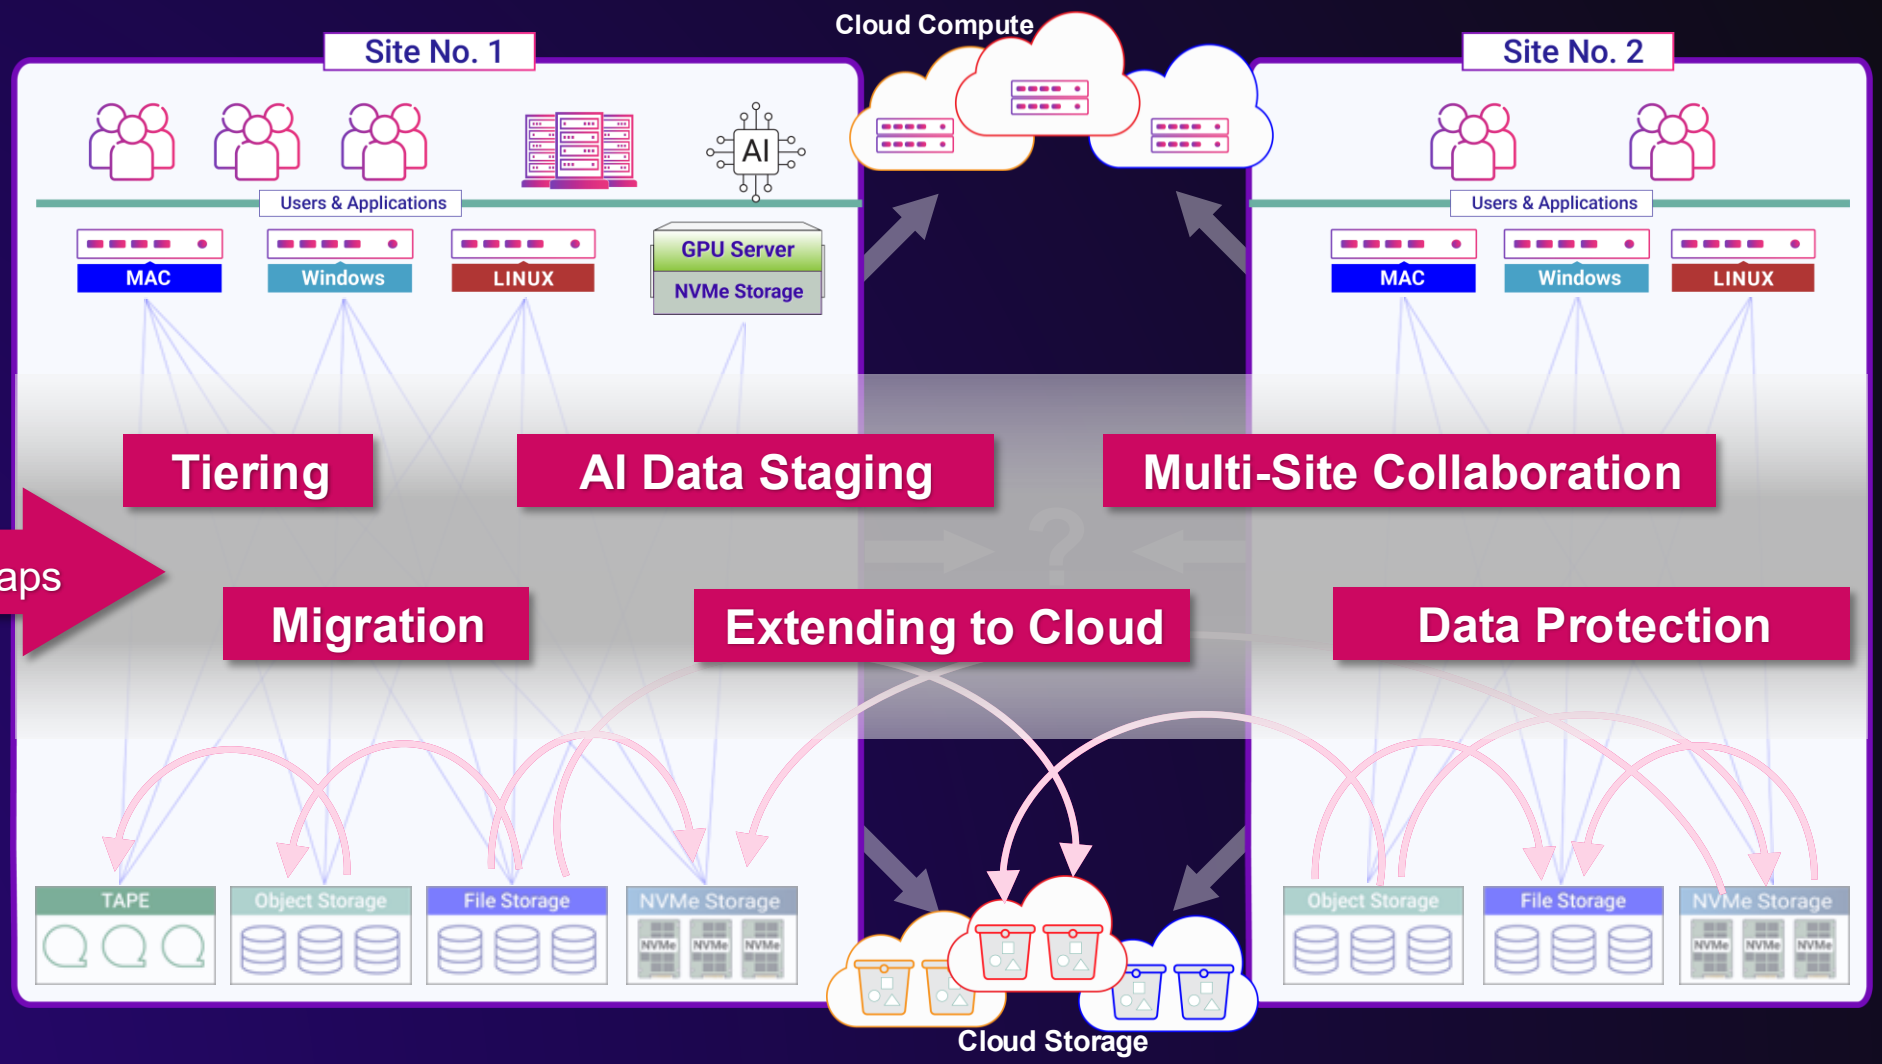

Multi-Vendor, Multi-Site, Multi-Cloud Storage

The Solution: Bridge Silos and Automate Data Placement

Unified High-Performance Access

- No proprietary client software needed
- Standard SMB, NFS & S3

Cross-platform Data Services

- Global data control with metadata intelligence

Accelerate Workflows

- High performance data movement

Any Storage

- Extreme performance Tier 0, or any third-party Storage

Non-Proprietary, Standards-based Access

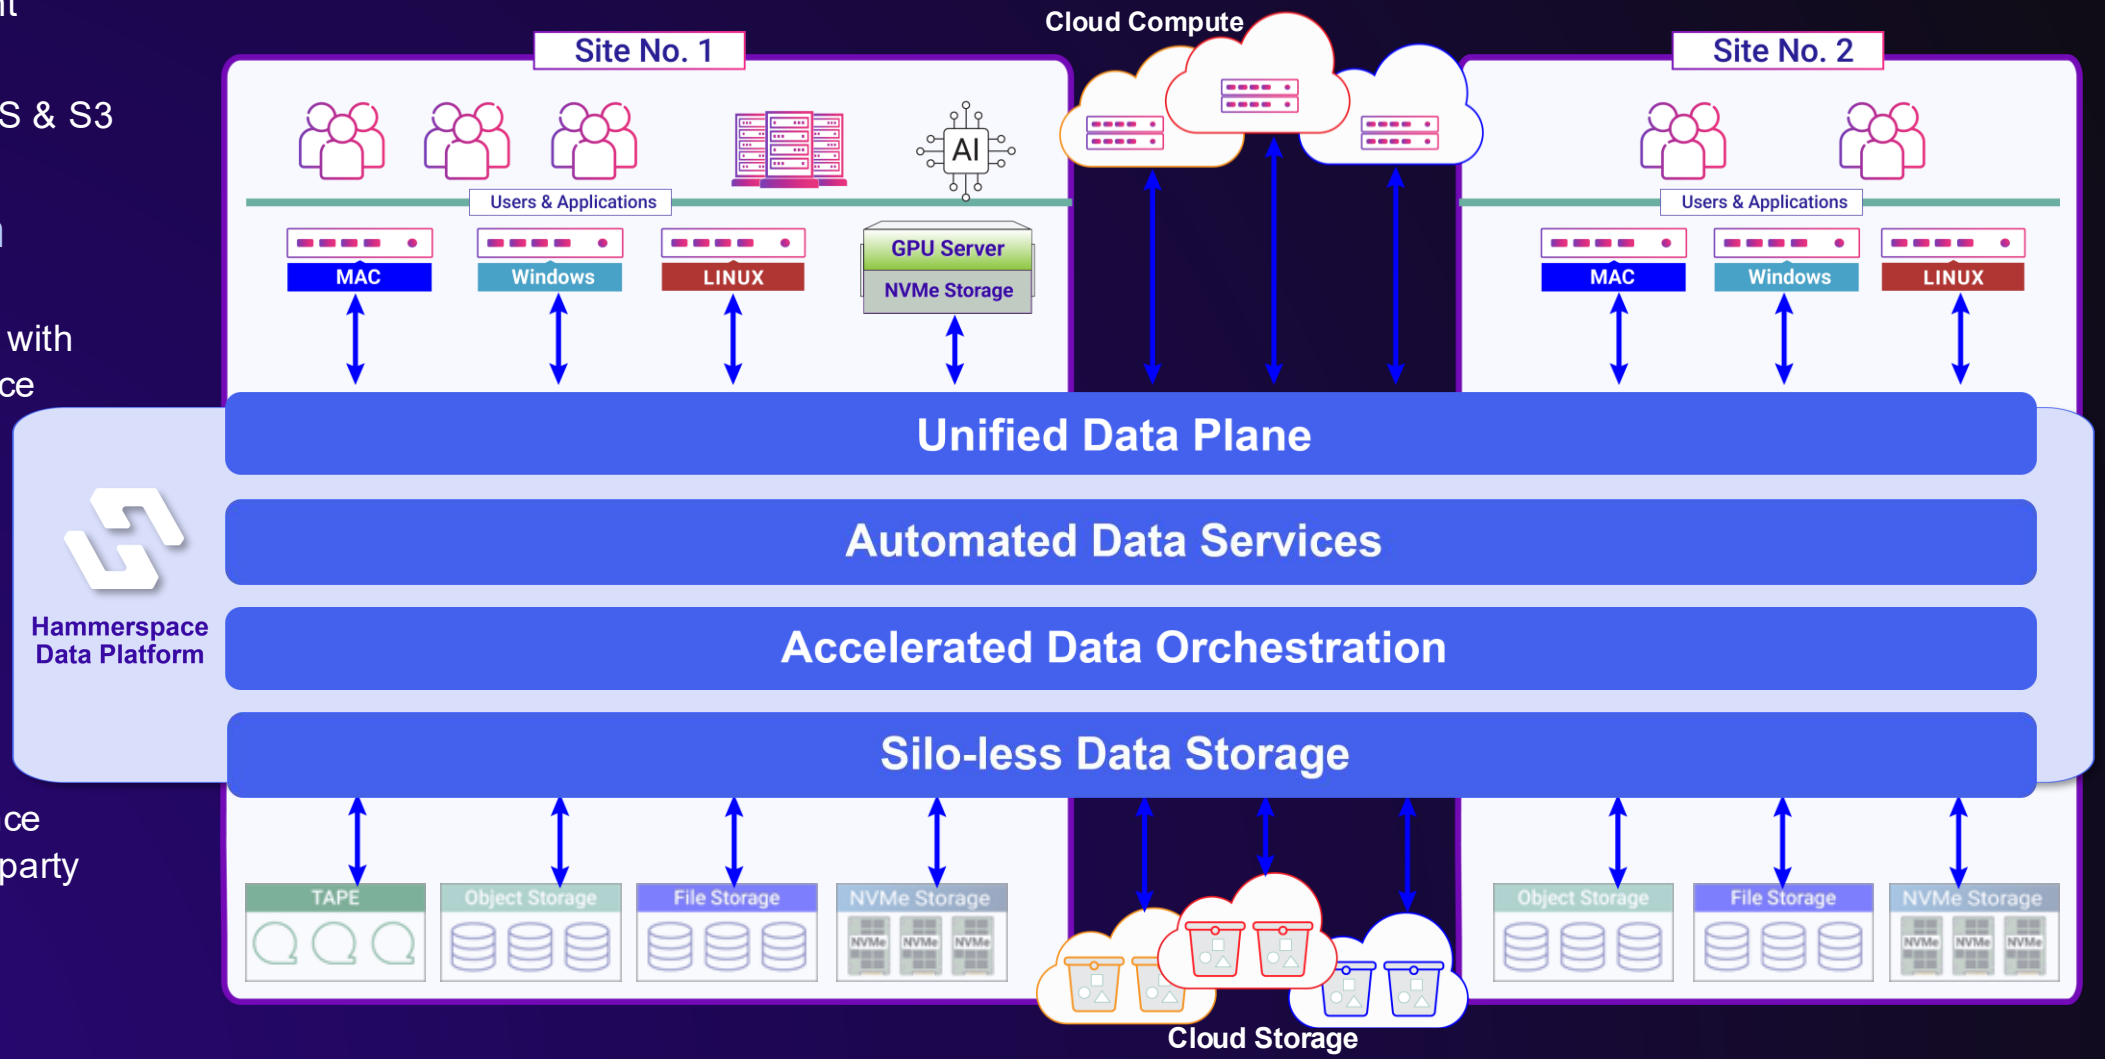

Multi-Vendor, Multi-Site, Multi-Cloud Storage

Experience Local Access to All Data -- Plus Global Control

Users simply access their files, as normal.

Admins control data and placement & protection without interrupting users

Users Have Seamless Access to All Data:

User View

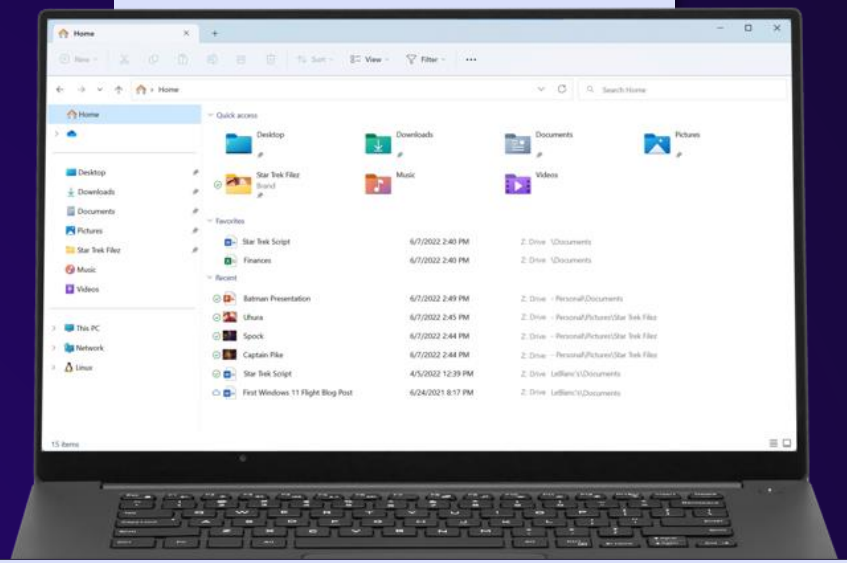

Users see data across all silos via standard protocols they are accustomed to.
No proprietary client software needed.

Admins Have Global Control of All Data Services:

Administrator View

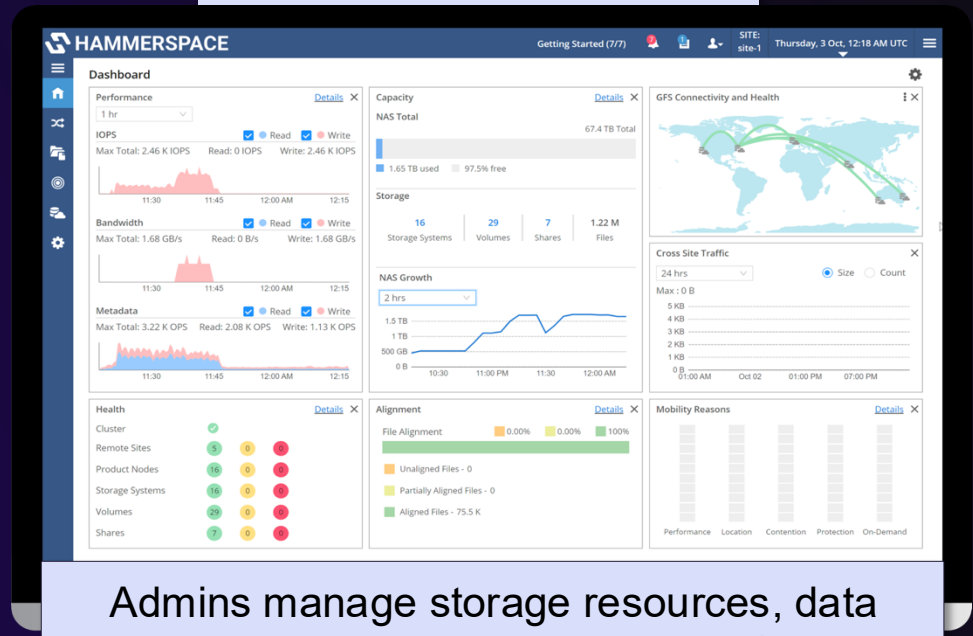

Admins manage storage resources, data placement, protections and workflows centrally across the entire data estate.

No Proprietary Client Software Needed!!

For Enterprise AI Adoption - Silos is Only the First Problem

Equally Important is Reducing Complexity

Speed-Bumps to Launching AI Projects

- Multiple Complex Steps
- Added Risk → Governance
- Huge upfront investment before ROI is known
- Requires Different Skillset than traditional IT

A Turn-key, Real-Time AI Data Pipeline for Enterprise

Taming Data Chaos

- Builds on Hammerspace core capabilities.
- Helps unlock their unstructured data for AI
- Removes steps for data collection, cataloging, and processing from data scattered across the enterprise
- Helps everyday enterprises start their AI journey by solving the biggest challenge to every project – AI Data readiness
- It's an agentic platform for AI Adoption

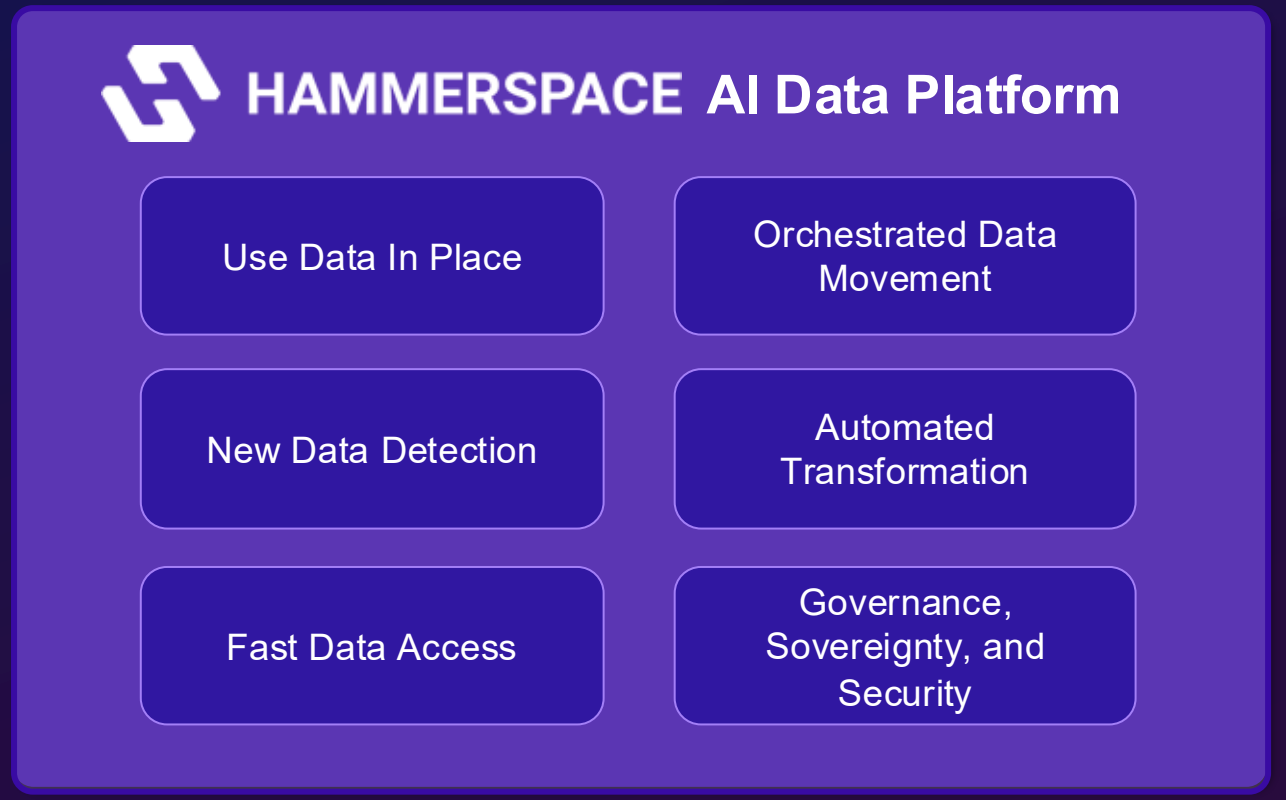

AI-Ready Data in Days, Not Months

The AI Data Platform for Accelerating Outcomes

Fast times to token are good, but time to VERY First token is essential

Traditional AI Solutions:

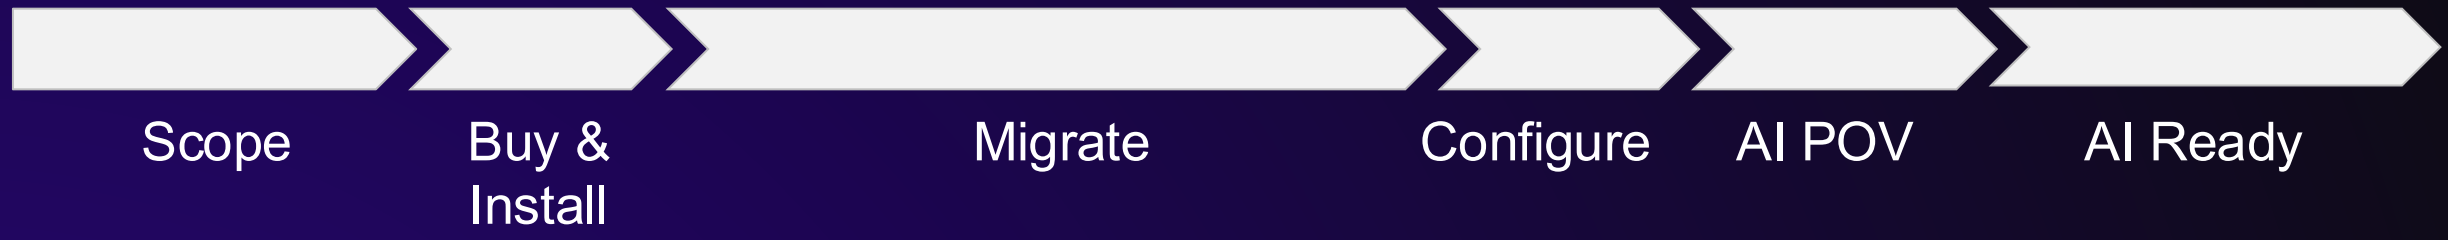

Hammerspace AIDP Solution:

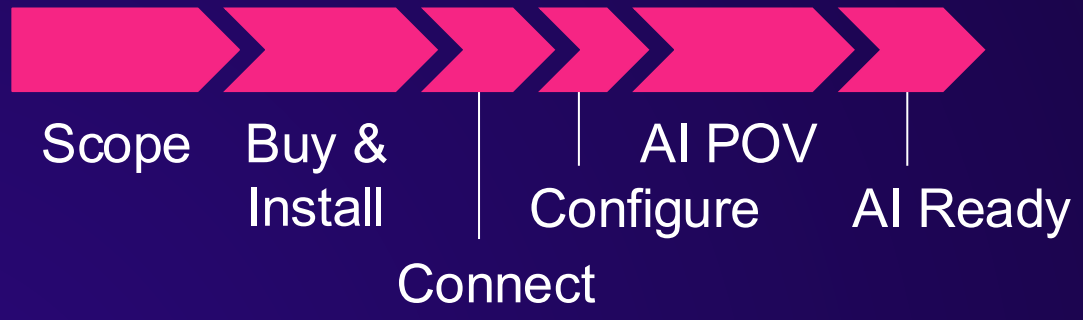

Over 10x Reduction in time spent compared to traditional AI Storage Solutions

Simplifies Transforms Existing Siloed Data into Unified AI-Ready Intelligence

Adds Key AI-Specific Features:

- Hammerspace MCP Server
- Vector DB & Data Intelligence Software
- Includes NVIDIA AI Enterprise Software

Built Upon Hammerspace Core Solution

Fully Integrated With Hardware:

- Cisco, Supermicro, Lenovo
- Supports Any Third-Party Hardware

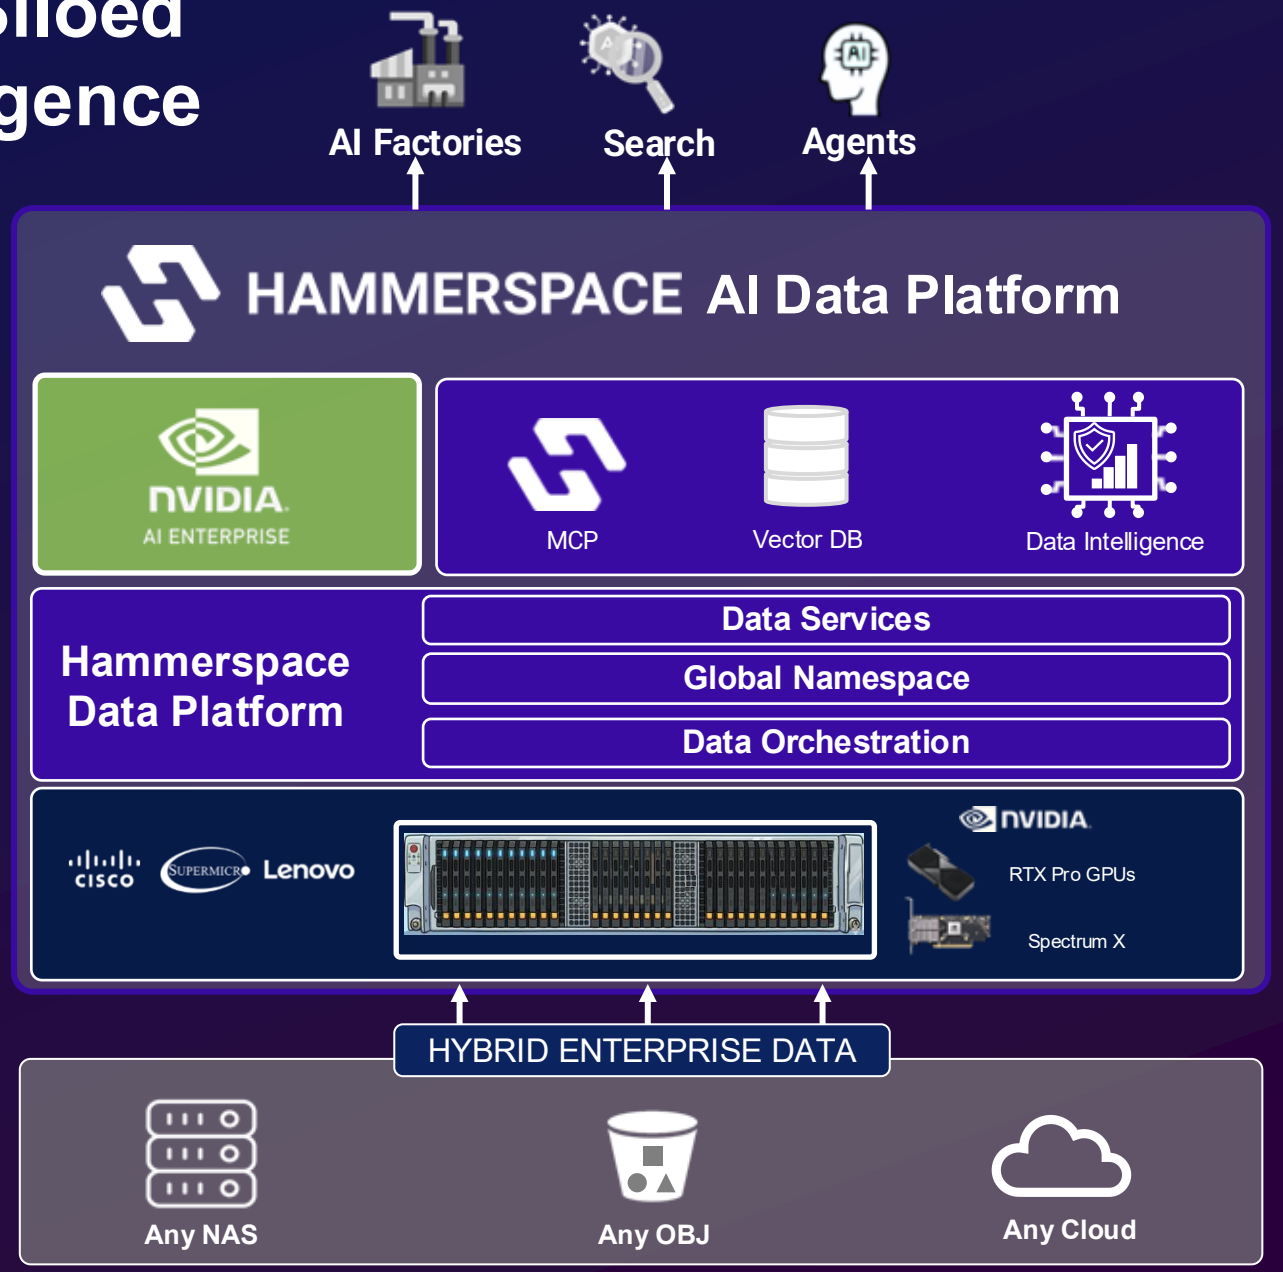

Hammerspace Data Platform with NVIDIA

RAG Pipeline with NVIDIA NIMS

NeMo Retriever Extraction → NeMo Retriever Embedding → NVIDIA cuVS Index → NeMo Retriever Reranking → VectorDB

MCP

NVIDIA Nemo Agent Toolkit

Validated Output

Consumption by AI Tools

MCP

Instant Embedding

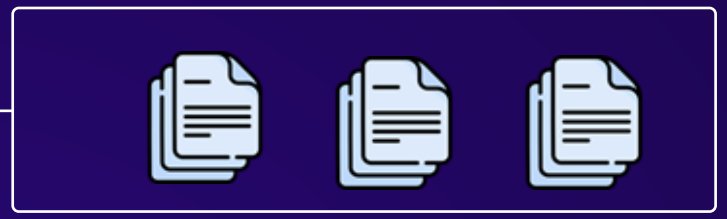

HAMMERSPACE

Data Platform

AI Agents, Factories & Automated Tools

AI Data Platform

Read Files in Place

Copy to us

Assimilate Existing Storage

Connect the Cloud

NVIDIA AI ENTERPRISE

Real-time Automated AI Processing

High Performance VectorDB

MCP Server

Fully integrated, Agentically Controlled

Data Intelligence & Governance

Enterprise Tools

AI Factory

Agents

HAMMERSPACE

AIDP Plain Language Query Interface

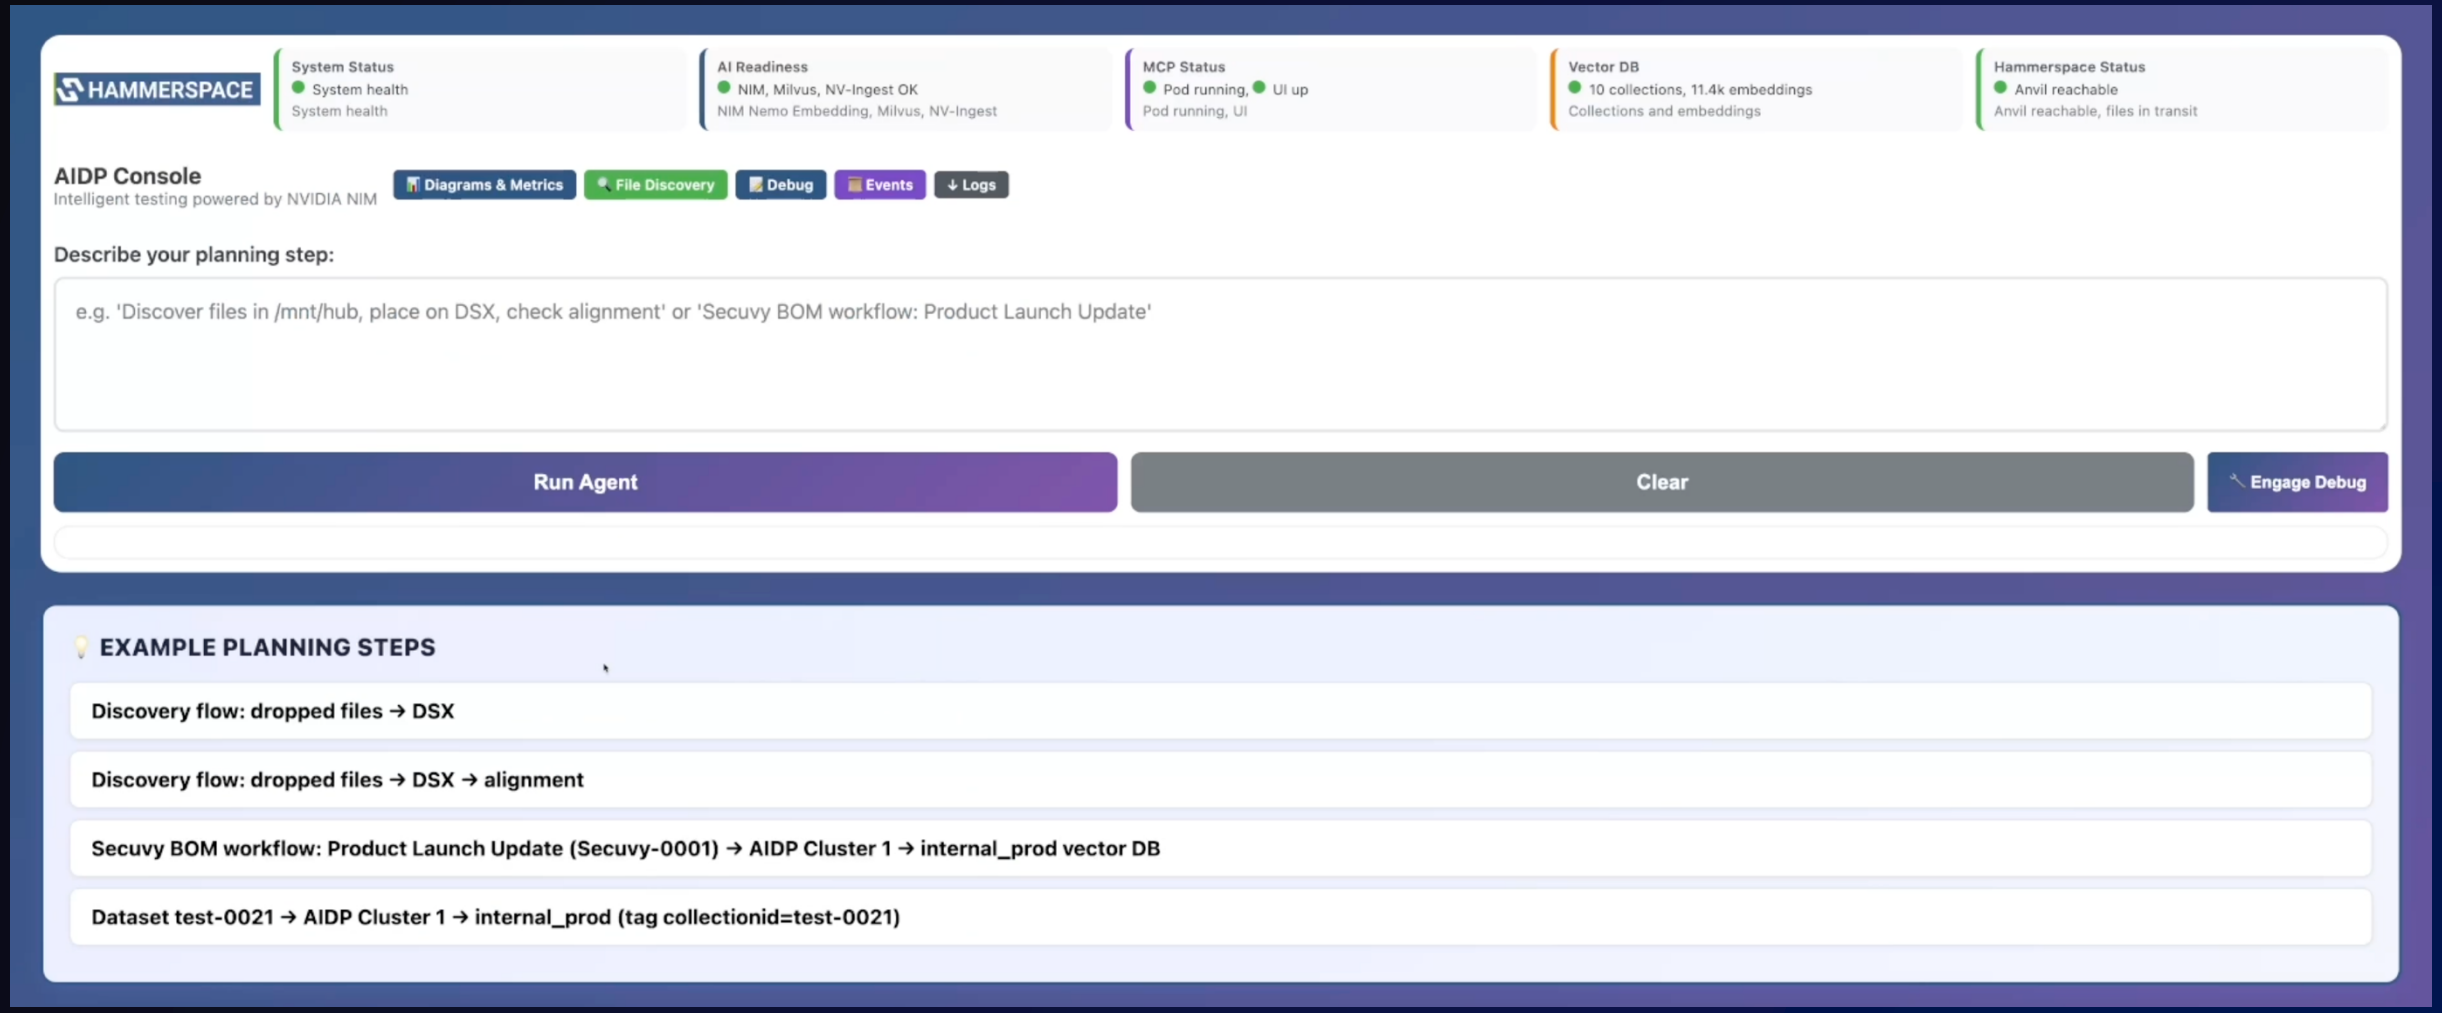

Rich Metadata Categorization

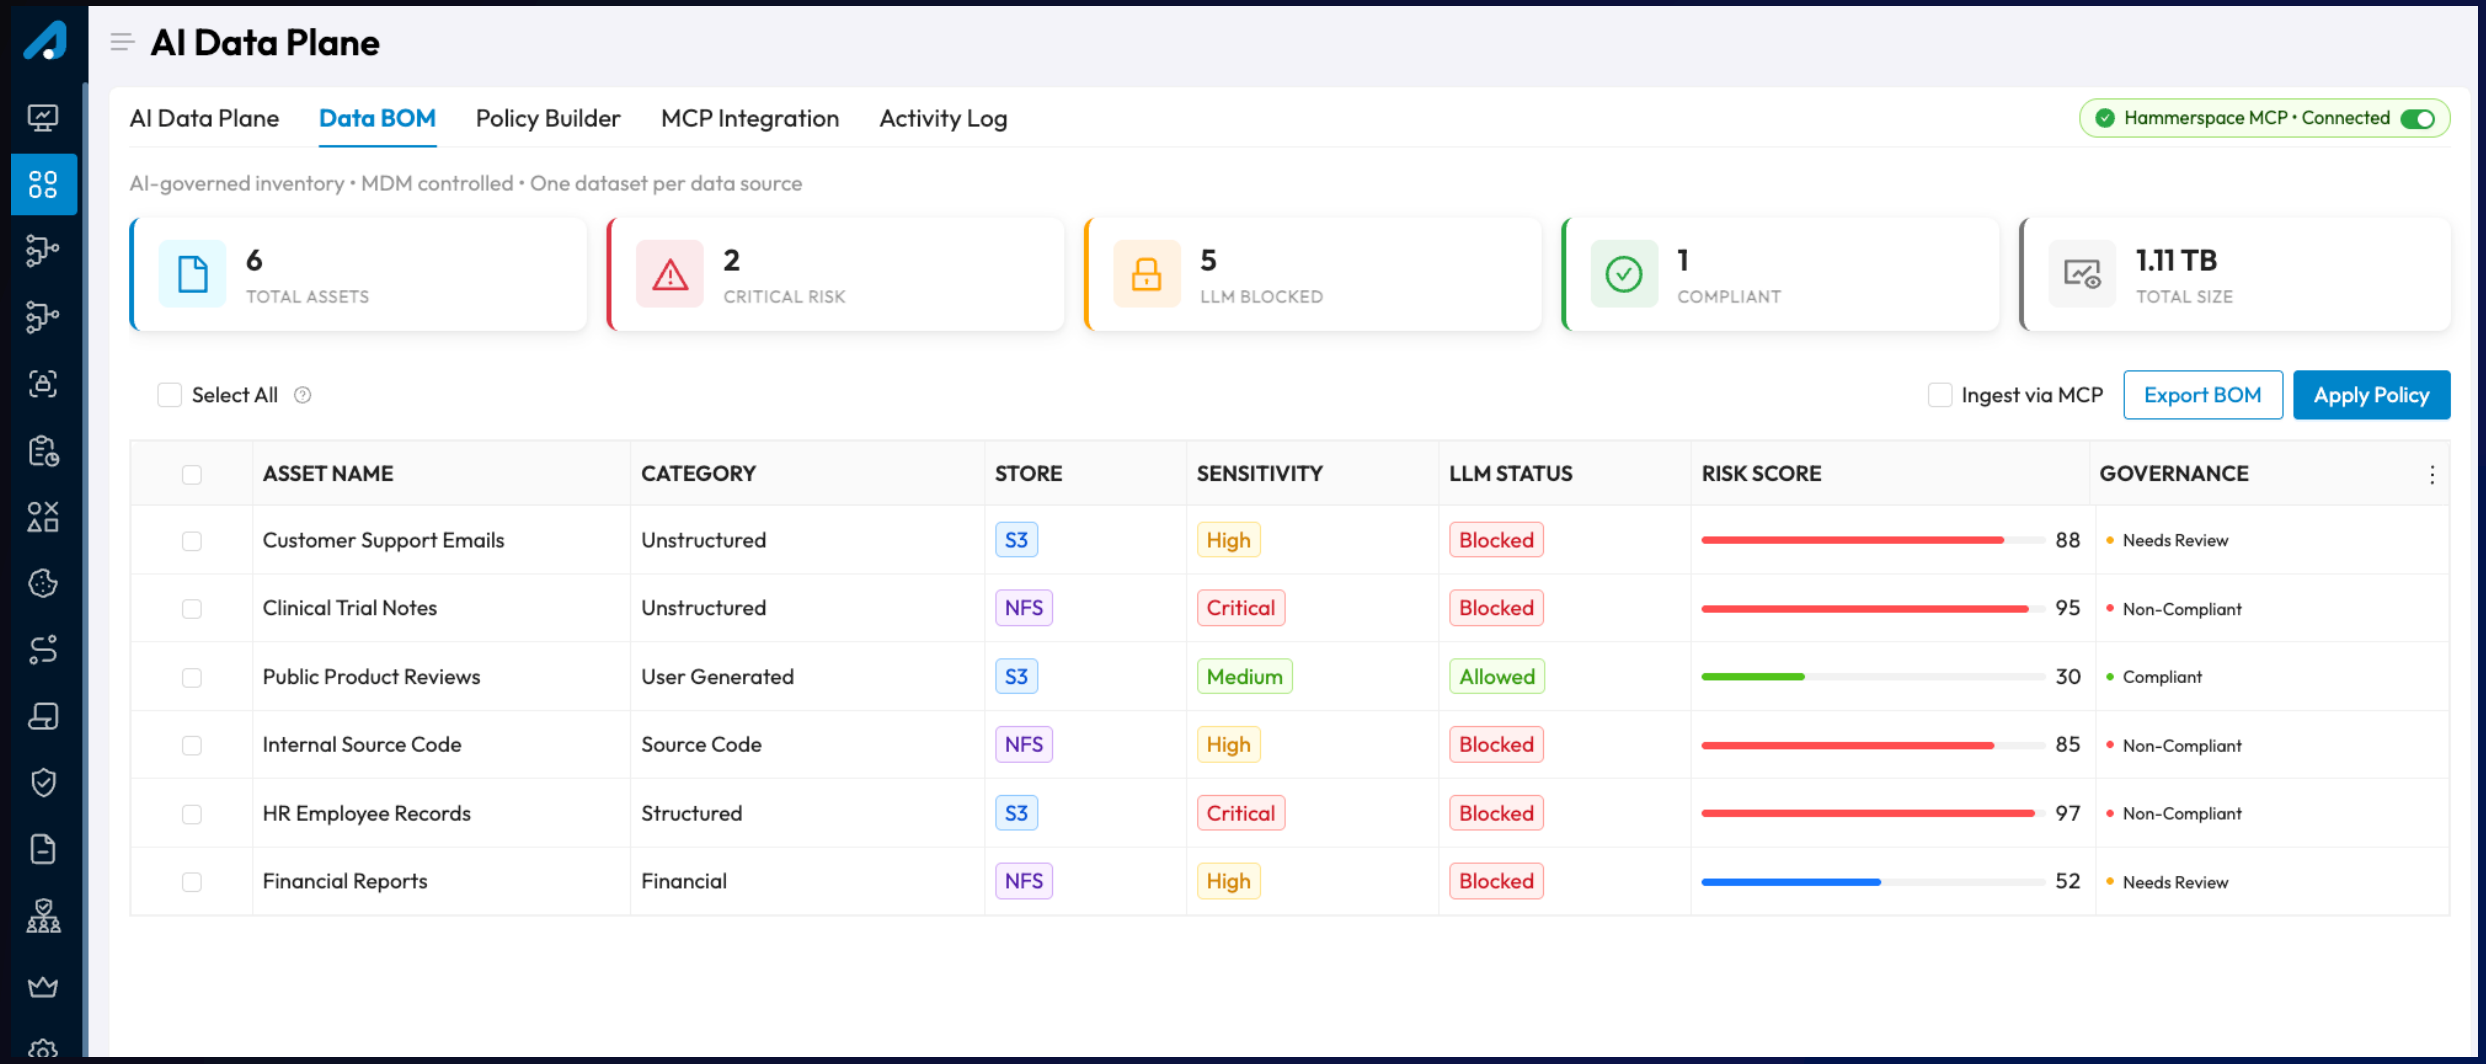

AI Data Platform

secuvyai

Rich Metadata Categorization

Files | Summary | Classification Details | **Business Tag Sub-Groups**

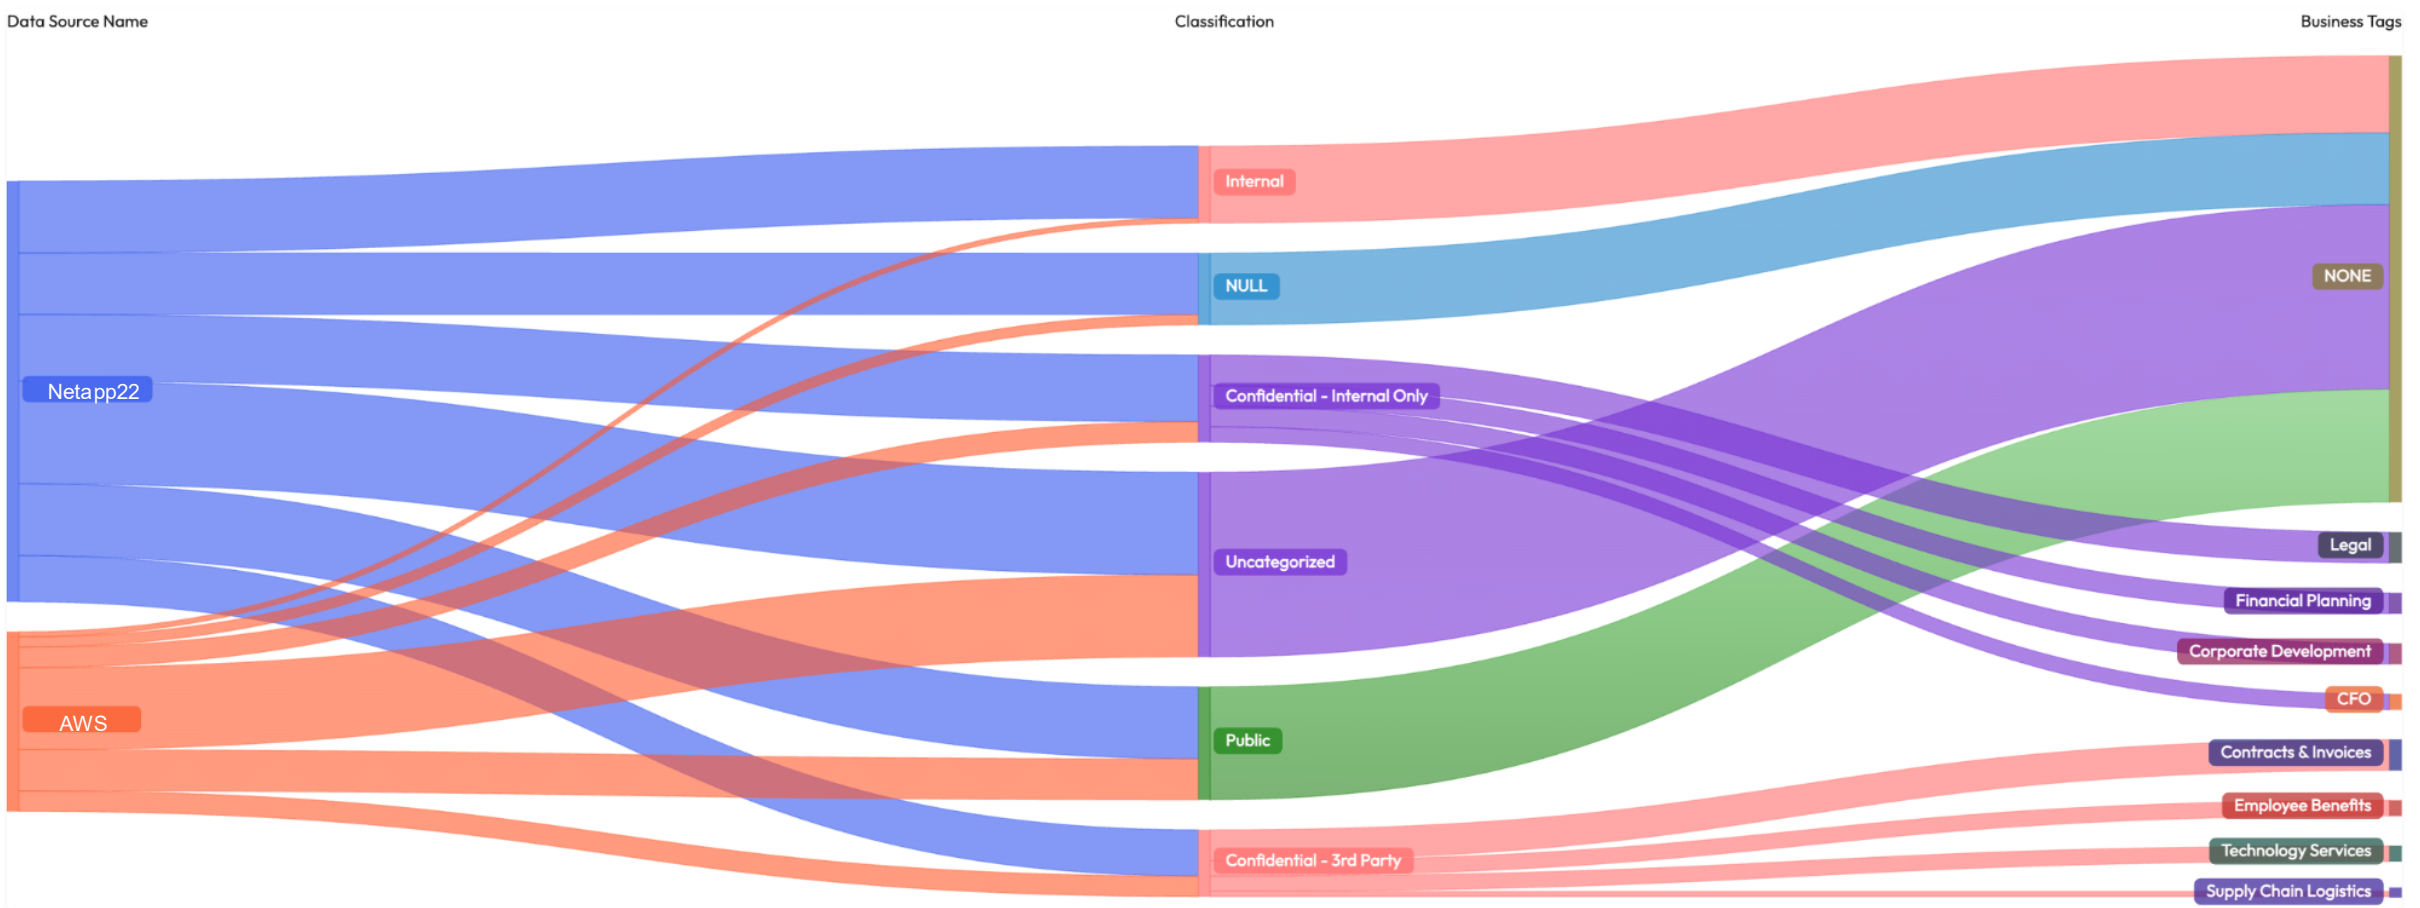

Automate Workloads with Intelligent Data Mobility

Orchestrate Data Pipelines to Reduce Operational Complexity

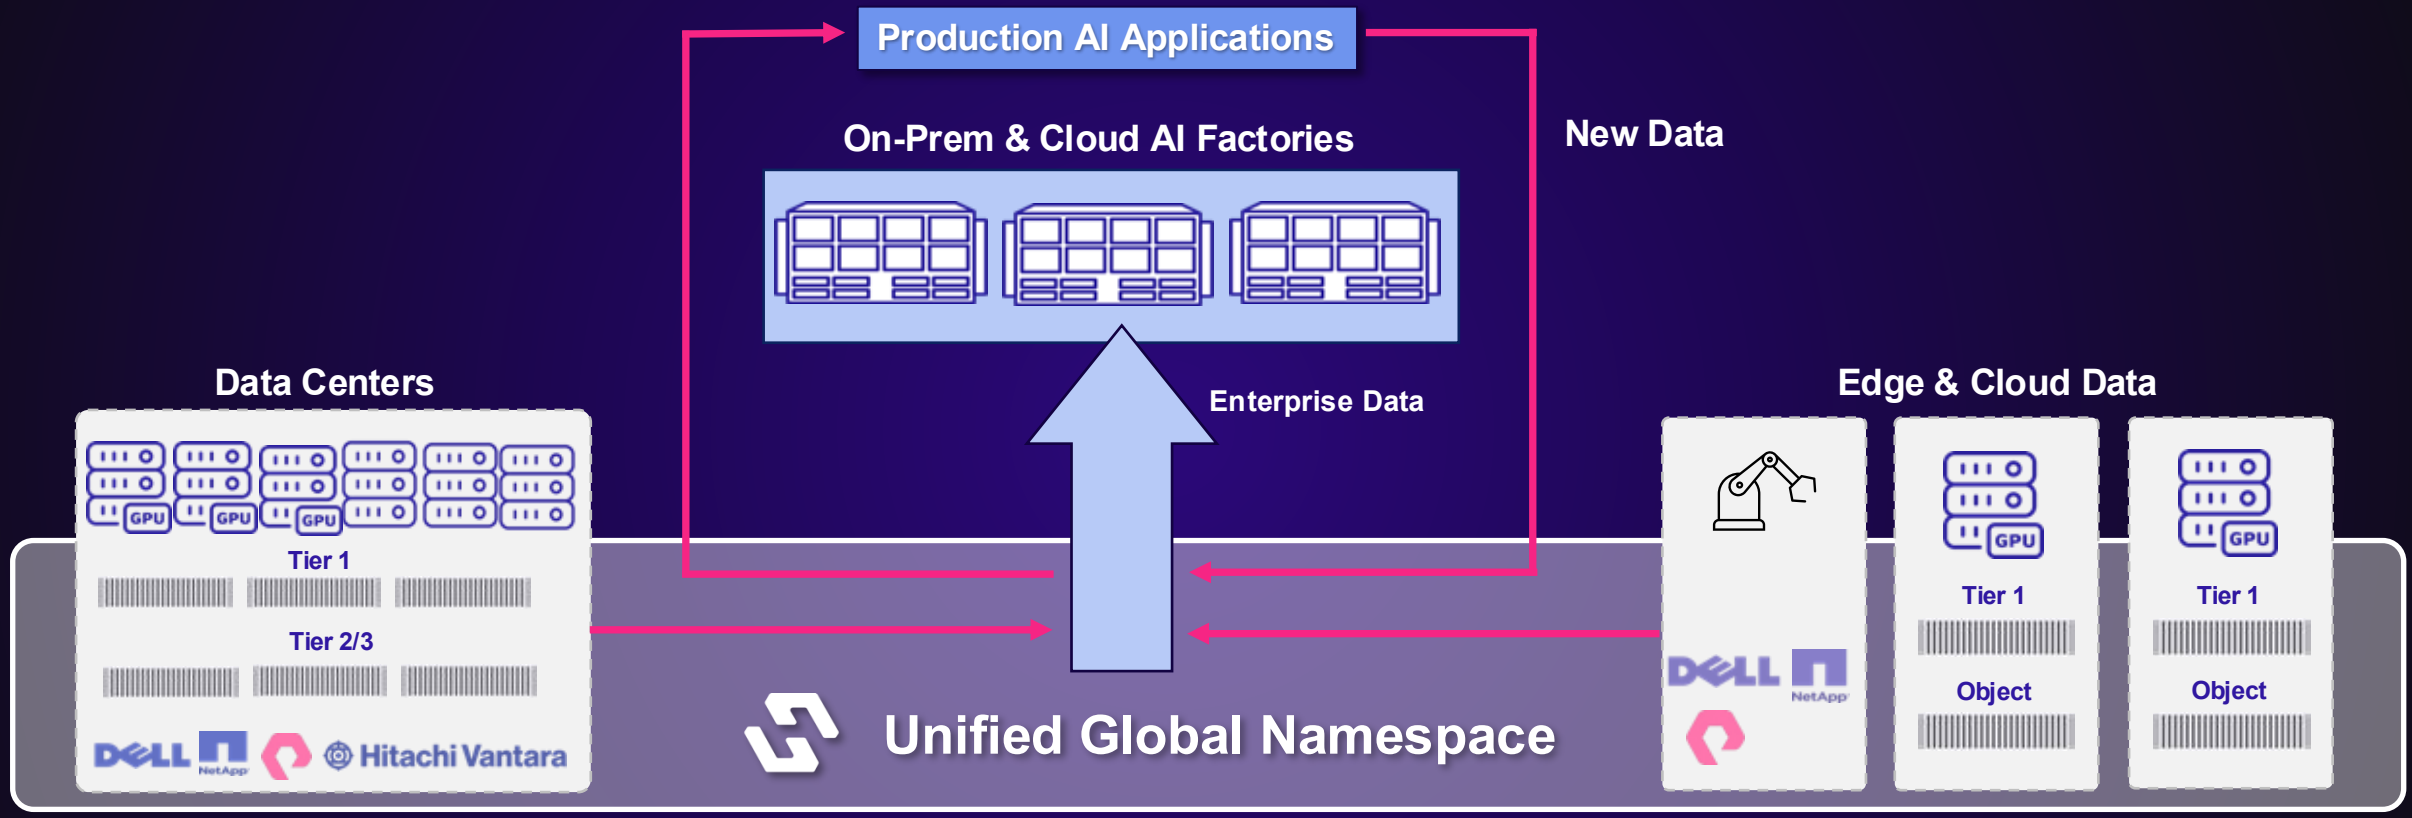

Intelligent, Non-Disruptive Data Movement

Move Data to the Right Tier/Location at the Right Time

Eliminate Data Gravity to Power AI anywhere

Hammerspace is the Data Fuel of AI

US GOVERNMENT

Department of DEFENSE

Department of ENERGY

AT-SCALE AI OPERATORS

Meta

zoom

NEOCLOUDS

US Government NEOCLOUD

Cirrascale®

Global File Access & Automated Data Orchestration

Non-Proprietary, Standards-based Access

REST API | pNFS with Flex Files | S3, NFS, SMB | MCP CSI

Hammerspace Data Platform

Unified Data Plane
High-Performance Parallel Global File System – Cross-Platform Global Namespace

Automated Data Services
Service-Level Objectives for Data Placement, Workflows, Data Protection, Custom Metadata Services

Accelerated Data Orchestration
Non-Disruptive Data Movement for High-Performance Uses, Tiering, Replication, Cloud, Multisite, etc.

Silo-less Data Storage
Tier 0, pNFSv4.2 with Flex Files, Optimized Object Storage, Cloud Storage

Multi-Vendor, Multi-Site, Multi-Cloud Storage

Thank You!

www.HAMMERSPACE.com